arXiv:1702.06099v1 [cs.DS] 20 Feb 2017

Preemptive Online Partitioning of Sequences *

Christian Konrad[1] and Tigran Tonoyan[1]

[1]Department of Computer Science and DIMAP, University of Warwick, UK,
c.konrad@warwick.ac.uk
[2]ICE-TCS, School of Computer Science, Reykjavik University, Iceland,
ttonoyan@gmail.com

February 21, 2017

Abstract

Online algorithms process their inputs piece by piece, taking irrevocable decisions for each data item. This model is too restrictive for most partitioning problems, since data that is yet to arrive may render it impossible to extend partial partitionings to the entire data set reasonably well.

In this work, we show that preemption might be a potential remedy. We consider the problem of partitioning online sequences, where $p-1$ separators need to be inserted into a sequence of integers that arrives online so as to create p contiguous partitions of similar weight. While without preemption no algorithm with non-trivial competitive ratio is possible, if preemption is allowed, i.e., inserted partition separators may be removed but not reinserted again, then we show that constant competitive algorithms can be obtained. Our contributions include:

We first give a simple deterministic 2-competitive preemptive algorithm for arbitrary p and arbitrary sequences. Our main contribution is the design of a highly non-trivial partitioning scheme, which, under some natural conditions and p being a power of two, allows us to improve the competitiveness to 1.68. We also show that the competitiveness of deterministic (randomized) algorithms is at least $\frac{4}{3}$ (resp. $\frac{6}{5}$).

For $p = 2$, the problem corresponds to the interesting special case of preemptively guessing the center of a weighted request sequence. While deterministic algorithms fail here, we provide a randomized 1.345-competitive algorithm for all-ones sequences and prove that this is optimal. For weighted sequences, we give a 1.628-competitive algorithm and a lower bound of 1.5.

1 Introduction

Online algorithms receive their inputs sequentially piece by piece. For each incoming piece of data (=request), the algorithm takes an immediate and irrevocable decision on how to process it. Taking good decisions can be challenging or even impossible, since decisions are based only on the requests seen and choices taken so far, and, in particular, they cannot be based on future requests. For many problems, taking a few bad decisions is forgivable and good competitive algorithms can nevertheless be designed (e.g. maximum matching [13], bin packing [24], k-server [17]), while for other problems, even a single bad decision may make it impossible to obtain non-trivial solutions (e.g. the maximum independent set problem [9]). Generally, a necessary condition for a problem to admit good online algorithms is that solutions can be *incrementally* built by extending partial solutions.

In this paper, we are interested in whether data partitioning problems can be solved online. In an online data partitioning problem, the input data X arrives online and is to be partitioned into p parts by computing a partitioning function $\phi : X \to \{1, 2, \ldots, p\}$ such that an application-specific cost function is optimized. Unfortunately, most partitioning problems are inherently non-incremental and thus poorly suited to the online model, which may explain why the literature on online data partitioning is exceptionally scarce (see related works section). This raises the following research questions:

*C. Konrad is supported by the Centre for Discrete Mathematics and its Applications (DIMAP) at Warwick University and by EPSRC award EP/N011163/1.

1. Is online data partitioning with provable guarantees really hopeless to achieve?
2. If so, how can we minimally augment the power of online algorithms to render data partitioning possible?

In this work, we address these questions with regards to the problem of *partitioning integer sequences*. It is one of the simplest data partitioning problems and thus a good candidate problem to answer the questions raised. We prove that even for this rather simple problem, non-trivial quality guarantees are indeed impossible to achieve in the online model. However, if we augment the online model with *preemption*, i.e., the ability to remove a previously inserted element from the solution, or, in the context of data partitioning, the ability to merge a subset of current partitions by removing previously inserted partition separators, then competitive ratios of at most 2 can be obtained.

Partitioning Online Integer Sequences. Let X be an integer sequence of length n, and $p \geq 2$ an integer. In the problem of partitioning integer sequences (abbreviated by PART), the goal is to partition X into p contiguous blocks (by determining the position of $p-1$ partition separators) such that the maximum weight of a block, denoted the bottleneck value of the partitioning, is minimized.

In the (non-preemptive) online model, parameter p is given to the algorithm beforehand, and the sequence X arrives online, integer by integer. When processing an integer, the algorithm has to decide whether or not to place a partition separator after the current integer. We are interested in the *competitive ratio* [2] of an algorithm, i.e., the (expected) ratio of the bottleneck value of the computed partitioning and the optimal bottleneck value of the input, maximized over all inputs. Placing no separator at all results in a single partition that is trivially p-competitive. We prove, however, that this is essentially best possible in the non-preemptive online model, even if the integer sequence is an all-ones sequence.

Given this strong impossibility result, we then augment the online model with preemption. When processing the current integer of the input sequence, a preemptive online algorithm for PART is allowed to remove a previously placed separator, which results in the merging of the two adjacent partitions incident to the separator. The removed separator can then be reinserted again (however, only after the current integer). In this paper, we show that the additional flexibility gained through preemption allows us to obtain algorithms with competitive ratio at most 2. Even though our original motivation for studying PART in the preemptive online model is the fact that non-trivial algorithms cannot be obtained in the non-preemptive case, preemptive online algorithms for PART are extremely space efficient (only the weights of partitions and positions of separators need to be remembered) and thus work well in a data streaming context for massive data sets.

Partitioning Continuous Online Flows. Algorithms for PART have to cope with the following two difficulties: First, the weights in the input sequence may vary hugely, which implies that algorithms cannot establish partitions of predictable weights. Second, algorithms need to find a way to continuously merge adjacent partitions while keeping the bottleneck value small. While the first point can be tackled via rounding approaches, the second constitutes the core difficulty of PART. We define a problem denoted FLOW, which abstracts away the varying weights of the integers and allows us focus on the second point. FLOW differs from PART in that the preemptive online algorithm is allowed to determine the weight of every incoming element (we now even allow for positive rational weights). The difficulty in FLOW stems from the fact that the algorithm is not aware of the total weight of the input. While at a first glance this problem appears to be substantially easier than PART, we show that any algorithm for FLOW can be used for PART while incurring an error term that depends on the ratio of the largest weight of an element and the total weight of the input sequence. FLOW can be seen as a continuous version of PART and can be interpreted as the problem of partitioning a continuous online flow (details follow in the preliminaries section).

Our Results. We first show that every algorithm for PART in the non-preemptive online model has an approximation ratio of $\Omega(p)$, even if the input is guaranteed to be a sequence of ones (**Theorem 1**). We then turn to the preemptive model and consider the special case $p = 2$ first, which corresponds to preemptively guessing the center of a weighted request sequence. It is easy to see that every deterministic algorithm for the $p = 2$ case has a competitive ratio of 2. We then give a randomized 1.345-competitive algorithm for unweighted sequences (**Theorem 2**) and prove that this is best possible (**Theorem 3**). We extend this algorithm to weighted sequences and give a 1.628-competitive algorithm (**Theorem 4**) and a lower bound of 1.5 on the competitiveness (**Theorem 5**).

For general p, we first give a simple deterministic 2-competitive online algorithm for PART (**Theorem 6**) and prove a lower bound of 4/3 (6/5) on the competitiveness of every deterministic (resp. randomized) algorithm (**Theorem 7**). We then turn to FLOW and give a highly non-trivial deterministic partitioning scheme with competitive ratio 1.68 (**Theorem 8**) for the case that p is a power of two, which constitutes the main contribution of this paper. This scheme translates to PART while incurring a small additive term in the competitive ratio that stems from the varying weights in PART. We discuss extensions of our scheme to arbitrary values of p and demonstrate experimentally that competitive ratios better than 2 can still be obtained. Last, we give a lower bound of 1.08 on the competitive ratio of every deterministic algorithm for FLOW (**Theorem 9**). Unlike the lower bounds for PART, this lower bound does not rely on the discrete properties of integers.

Techniques. Consider first the $p = 2$ case and all-ones sequences, which corresponds to preemptively guessing the center of the request sequence. One potential technique is *reservoir sampling* [25], which allows the sampling of a uniform random element while processing the input sequence. It naturally suits the preemptive online model and can be used to place a separator at a uniform random position in the request sequence, giving a randomized algorithm with expected competitive ratio 1.5. We show that the *randomized geometric guessing* technique allows us to improve on this bound: For a random $\delta \in (0, 1)$ and a carefully chosen value X, reset the single separator to the current position every time the total weight seen so far equals $\lceil X^{\delta \cdot i} \rceil$, for $i \in \{1, 2, \dots\}$, giving a 1.345-competitive algorithm. Via Yao's principle, we prove that this algorithm is optimal. We then analyze essentially the same algorithm on weighted sequences and show that it is 1.632-competitive.

For general p, consider the following algorithm for all-ones sequences of unknown length n (assume also that p is even): First, fill all partitions with weight $w = 1$. Whenever partitions are entirely filled, merge them pairwise creating $p/2$ partitions of weight $2w$ and then update $w \leftarrow 2w$. Then w always constitutes the bottleneck value of this partitioning. After the merging, fill the empty partitions with weight w and repeat. Note that the optimal bottleneck value *opt* is bounded as $opt \geq \frac{n}{p}$. Since at least half of all partitions computed by the algorithm have weight w, it holds that $w\frac{p}{2} \leq n$, which together with $opt \geq \frac{n}{p}$ implies $w \leq 2opt$, giving a 2-competitive algorithm. The 2-competitive algorithm given in this paper is based on the intuition provided and also works for sequences with arbitrary weights.

To go beyond the competitiveness of 2, consider the key moment that leads to the 2-competitiveness of the above algorithm: Just after merging all p partitions of weight w into $p/2$ partitions of weight $2w$, the competitive ratio is 2. To avoid this situation, note that merging only a single pair of partitions would not help. Instead, after merging two partitions, we need to guarantee that the new bottleneck value is substantially smaller than twice the current bottleneck value. This implies that the weights of the merged partitions cannot both be close to the current bottleneck value, and it is hence beneficial to establish and merge partitions with different weights. In fact, our 1.08 lower bound for FLOW makes use of this observation: If at some moment the competitive ratio is too good, then most partitions have similar weight, which implies that when these partitions have to be merged in the future the competitive ratio will be large.

From an upper bound perspective, a clever merging scheme is thus required, which establishes partitions of different weights and, in particular, remains analyzable. We give such a scheme for FLOW, when p is a power of two. As an illustration, consider the case $p = 4$ as depicted in Table 1. Recall that in FLOW we are allowed to determine the weight of the incoming elements. We first initialize the partitions with values x, x^2, x^3, x^4, for $x = 2^{\frac{1}{4}}$, and evolve the partitions as in the table. Note that, at every moment, all partition weights are different, but at the same time never differ by more than a factor of 2. A key property is that at the end of the scheme the weights of the partitions are a multiple of the initial weights of the partitions. This allows us to repeate the scheme and limits the analysis to one cycle of the scheme. We prove that the competitive ratio of our scheme is 1.68 for every p that is a power of two. We discuss how our scheme can be extended to other values of p and demonstrate experimentally that competitive ratios better than 2 can still be obtained.

Further Related Work. The study of PART in the offline setting has a rich history with early works dating back to the 80s [1, 7, 11, 18, 19, 14, 10, 20, 21]. After a series of improvements, Frederickson gave a highly non-trivial linear time algorithm [7]. PART finds many applications, especially in load balancing scenarios (e.g. [20, 15, 16]). It has recently been studied in the context of streaming algorithms where it serves as a building block for partitioning XML documents [16].

Recently, Stanton and Kliot [23] expressed interest in simple online strategies for data partitioning.

P_1	P_2	P_3	P_4	next
$\boldsymbol{x}$	$\boldsymbol{x^2}$	x^3	x^4	x^2
$x(1+x)$	x^3	x^4	$\boldsymbol{x^2}$	$\boldsymbol{x^3}$
$x(1+x)$	$\boldsymbol{x^3}$	$\boldsymbol{x^4}$	$x^2(1+x)$	x^4
$x(1+x)$	$x^3(1+x)$	$x^2(1+x)$	$\boldsymbol{x^4}$	$\boldsymbol{x^5}$
$\boldsymbol{x(1+x)}$	$\boldsymbol{x^3(1+x)}$	$x^2(1+x)$	$x^4(1+x)$	$x^3(1+x)$
$x(1+x)(1+x^2)$	$\boldsymbol{x^2(1+x)}$	$\boldsymbol{x^4(1+x)}$	$x^3(1+x)$	$x^5(1+x)$
$x(1+x)(1+x^2)$	$x^2(1+x)(1+x^2)$	$\boldsymbol{x^3(1+x)}$	$\boldsymbol{x^5(1+x)}$	$x^4(1+x)$
$x(1+x)(1+x^2)$	$x^2(1+x)(1+x^2)$	$x^3(1+x)(1+x^2)$	$\boldsymbol{x^4(1+x)}$	$\boldsymbol{x^6(1+x)}$
$x(1+x)(1+x^2)$	$x^2(1+x)(1+x^2)$	$x^3(1+x)(1+x^2)$	$x^4(1+x)(1+x^2)$	

Table 1: A partitioning scheme for Flow for $p = 4$. The bold elements are merged next.

They studied online graph partitioning heuristics[1] for the balanced graph partitioning problem and demonstrated experimentally that simple heuristics work well in practice. Stanton later analyzed the behavior of these heuristics on random graphs and gave good quality bounds [22]. Interestingly, besides this line of research, we are unaware of any further attempts at online data partitioning.

Many works provide additional power to the online algorithm. Besides preemption, common resource augmentations include lookahead (e.g. [8]), distributions on the input (e.g. [12]), or advice (e.g. [4]). Preemptive online algorithms have been studied for various online problems. One example with a rich history is the matching problem (e.g. [6, 27, 5, 3]).

Outline. In Section 2, we formally define the studied problems and the preemptive online model. In Section 3, we prove that every non-preemptive algorithm for Part has a competitive ratio of $\Omega(p)$. Then, we give our algorithms and lower bounds for Part for the special case $p = 2$ in Section 4. All our results for Part and Flow for arbitrary p are given in Section 5.

2 Preliminaries

Partitioning Integer Sequences. In this paper, we study the following problem:

Definition 1 (Partitioning Integer Sequences)**.** Let $X = w_1, \ldots, w_n \in \mathbb{N}^n$ be an integer sequence, and let $p \in \mathbb{N}$ be an integer. The problem of partitioning integer sequences consists of finding separators $S = s_1, \ldots, s_{p-1}$ such that $1 = s_0 \leq s_1 \leq s_2 \leq \cdots \leq s_{p-1} \leq s_p = n + 1$ and the maximum weight of a partition is minimized, i.e., $\max_{j \in \{0,\ldots,p-1\}} \sum_{i=s_j}^{s_{j+1}-1} w_i$ is minimized. The weight of a heaviest partition is the *bottleneck value* of the partitioning. This problem is abbreviated by Part.

Online Model. In the online model, parameter p is given to the algorithm beforehand, and the integers $X = w_1, \ldots, w_n$ arrive online. Upon reception of an integer w_i (also called a *request*), the algorithm has to decide whether to place a partition separator after w_i. In the non-preemptive online model, placing a separator is final, while in the preemptive model, when processing w_i previously placed separators may be removed. The total number of separators in place never exceeds $p - 1$ and separators can only be inserted at the current request.

Competitive Ratio. The *competitive ratio* of a deterministic online algorithm for Part is the ratio between the bottleneck value of the computed partitioning and the bottleneck value of an optimal partitioning, maximized over all potential inputs. If the algorithm is randomized, then we are interested in the *expected competitive ratio*, where the expectation is taken over the coin flips of the algorithm.

Partitioning Online Flows. We connect Part to the problem of partitioning a continuous online flow, abbreviated by Flow.

Definition 2 (Preemptive Partitioning of a Continuous Online Flow)**.** In Flow, time is continuous starting at time 0. Flow enters the system with unit and constant speed such that at time t, the total volume of flow $\int_0^t 1\,dx = t$ has been injected. The goal is to ensure that when the flow stops at time

[1] Phrased in the context of streaming algorithms, but their algorithms are in fact online

$t_{max} \geq t_0$ (t_0 is an arbitrarily small initial warm-up period), which is unknown to the algorithm, the total amount of flow is partitioned into p parts such that the weight of a heaviest partition is minimized. More formally, at time $t_{max} \geq t_0$, the objective is that partition separators $S = s_1, \ldots, s_{p-1}$ with $0 = s_0 \leq s_1 \leq \cdots \leq s_{p-1} \leq s_p = t$ are in place such that: $\max_{j \in \{0,\ldots,p-1\}} \int_{s_j}^{s_{j+1}} 1 \, dt = \max_{j \in \{0,\ldots,p-1\}} s_{j+1} - s_j$ is minimized. Similar to PART, in the preemptive online model, partition separators can only be inserted at the current time t, and previously inserted partition separators can be removed.

Even though FLOW is defined as a continuous problem, it can be seen as a special case of PART where the algorithm can determine the weight of every incoming element.

3 An $\Omega(p)$ Lower Bound In The Non-preemptive Online Model

Solving PART in the non-preemptive online model is difficult since the total weight of the input sequence is unknown to the algorithm. Since inserted partition separators cannot be removed, any partition created by the algorithm may be too small if the input sequence is heavier than expected. This intuition is formalized in the following theorem:

Theorem 1. For every $p \geq 2$, every randomized non-preemptive online algorithm for PART has expected competitive ratio $\Omega(p)$, even on all-ones sequences.

Proof. Let $k = p^2$ and let $\Sigma = \{\sigma_1, \ldots, \sigma_k\}$ be a set of request sequences where σ_i is the all-ones sequence of length $4ip$. Let $\mathcal{A}$ be a randomized algorithm for PART, and assume that its expected competitive ratio on every instance of Σ is at most c. Then by Yao's lemma, there exists a deterministic algorithm $\mathcal{A}_d$ with average approximation ratio at most c over the instances of Σ.

Let $S = \{s_1, \ldots, s_{p-1}\}$ denote the set of separators output by $\mathcal{A}_d$ on σ_k. Note that since σ_i is a prefix of σ_k, the output of $\mathcal{A}_d$ on σ_i is a subset of the separators S. Partition now Σ into Σ_0 and Σ_1 such that the approximation ratio of $\mathcal{A}_d$ on $\sigma \in \Sigma_1$ is strictly smaller than $p/2$, and the approximation ratio of $\mathcal{A}_d$ on $\sigma \in \Sigma_0$ is at least $p/2$. Consider now a $\sigma_i \in \Sigma_1$. Then, there exists a separator $s(i) \in S$ with $1/4 \cdot 4ip \leq s_j \leq (3/4) \cdot 4ip$, since otherwise the bottleneck value of the partitioning output by $\mathcal{A}_d$ on σ_i was at least $2pi$. This in turn would imply that the approximation ratio was at least $2pi/(4i) = \frac{1}{2}p$ (since the optimal bottleneck value on σ_i is $4i$), contradicting the fact that $\sigma \in \Sigma_1$. Next, note that the separators $s(i)$ and $s(j)$, for $i \neq j$ and $\sigma_i, \sigma_j \in \Sigma_1$, are necessarily different. Thus, since the number of separators is $p - 1$, the size of Σ_1 is bounded by $p - 1$. The average approximation factor c of $\mathcal{A}_d$ on instances S is thus at least

$$c \geq \frac{|\Sigma_0| \cdot \frac{p}{2} + |\Sigma_1| \cdot 1}{|\Sigma|} = \frac{(k - (p-1))\frac{p}{2}}{k} > \frac{p}{2} - \frac{p^2}{2k} \geq \frac{p}{2} - \frac{p^2}{2k} = \frac{p}{2} - \frac{1}{2} \, .$$

□

4 Guessing the Center: Part for $p = 2$

In this section, we consider PART for $p = 2$, i.e., a single separator needs to be introduced into the request sequence so as to split it into two parts of similar weight. We start with all-ones sequences and present an asymptotically optimal preemptive online algorithm. Then, we show how to extend this algorithm to sequences of arbitrary weights.

4.1 All-ones Sequences

The special case $p = 2$ on all-ones sequences corresponds to preemptively guessing the center of a request sequence of unknown length. Deterministic algorithms cannot achieve a competitive ratio better than 2 here, since request sequences that end just after a deterministic algorithm placed a separator give a competitive ratio of 2.

Observation 1. Every deterministic preemptive online algorithm for PART with $p = 2$ has a competitive ratio of 2.

Using a single random bit, the competitiveness can be improved to 1.5. This *barely-random* algorithm is given in Appendix A. Using $O(\log n)$ random bits, we can improve the competitive ratio to 1.344, which is best possible, and will be presented now.

Algorithm 2 Alg. $\mathcal{A}_x$ for PART and $p = 2$

```
Choose uniform random δ ∈ (0,1), i → 0
for each request j = 1,...,n do
  if j = ⌈x^{i+δ}⌉ then
    move separator to current position
    i → i + 1
  end if
end for
```

Algorithm $\mathcal{A}_x$, as depicted in Algorithm 2, is parametrized by a real $x > 2$, which will be optimized in Theorem 2. It moves the separator to the current position as soon as the $\lceil x^{i+\delta} \rceil$-th request is processed, where i is any integer and $\delta \in (0, 1)$ is a random number.

Remark. The continuous random variable δ is only taken for convenience in the analysis; a bit precision of $O(\log n)$ is enough.

In the following, denote by $R^n_{\mathcal{A}_x}$ the competitive ratio of $\mathcal{A}_x$ on a sequence of length n.

Theorem 2. There is a constant $x \approx 3.052$ such that $\mathbb{E}[R^n_{\mathcal{A}_x}] \approx 1.344 + O(n^{-1})$.

Proof. Let $\alpha \in [0, 1)$ and $i \in \mathbb{N}$ be such that $n = 2x^{i+\alpha}$. Then, the bottleneck value of an optimal partition is $\lceil \frac{n}{2} \rceil = \lceil x^{i+\alpha} \rceil$. We now bound the bottleneck value of the computed partitioning $r^n_{\mathcal{A}_x}$, which depends on various ranges of α and δ. We distinguish two ranges for α, and within each case, we distinguish three ranges of δ:

Case 1: $\alpha > 1 - \log_x 2$ (note that we assumed $x > 2$). In order to bound $r^n_{\mathcal{A}_x}$, we split the possible values of δ into three subsets:

- If $\delta \in (0, \alpha + \log_x 2 - 1]$, then we have that $x^{\delta+i+1} \le 2x^{i+\alpha} = n$. In this case, the bottleneck value is $r^n_{\mathcal{A}_x} = x^{\delta+i+1} + O(1)$.
- If $\delta \in (\alpha + \log_x 2 - 1, \alpha]$, then we have that $x^{\delta+i+1} > n$ but $x^{\delta+i} \le \frac{n}{2}$. In this case, $r^n_{\mathcal{A}_x} = n - x^{\delta+i} + O(1) = 2x^{\alpha+i} - x^{\delta+i} + O(1)$.
- If $\delta \in (\alpha, 1)$, then we have that $x^{\delta+i+1} > n$ and $x^{\delta+i} \in (\frac{n}{2}, n)$. In this case, $r^n_{\mathcal{A}_x} = x^{\delta+i}$.

Using these observations, we can bound the expected competitive ratio as follows:

$$\begin{aligned}
\mathbb{E}[R^n_{\mathcal{A}_x}] &= \int_0^{\alpha+\log_x 2-1} \frac{x^{\delta+i+1}}{x^{i+\alpha}} d\delta + \int_{\alpha+\log_x 2-1}^{\alpha} \frac{2x^{\alpha+i} - x^{\delta+i}}{x^{i+\alpha}} d\delta + \int_{\alpha}^{1} \frac{x^{\delta+i}}{x^{i+\alpha}} d\delta + O(n^{-1}) \\
&= \frac{1}{x^\alpha} \left(\int_0^{\alpha+\log_x 2-1} x^{\delta+1} d\delta + \int_{\alpha+\log_x 2-1}^{\alpha} (2x^\alpha - x^\delta) d\delta + \int_{\alpha}^{1} x^\delta d\delta \right) + O(n^{-1}) \\
&= 2 - 2\log_x 2 + \frac{2}{x \ln x} + O(n^{-1}).
\end{aligned}$$

Case 2: $\alpha \le 1 - \log_x 2$. We deal with this case similarly, but we need to group the possible values for δ in a different way:

- If $\delta \in (0, \alpha]$, then $x^{\delta+i+1} > n$ but $x^{\delta+i} \le \frac{n}{2}$. In this case, $r^n_{\mathcal{A}_x} = n - x^{\delta+i} + O(1)$.
- If $\delta \in (\alpha, \alpha + \log_x 2]$, then $x^{\delta+i} > \frac{n}{2}$ and $x^{\delta+i} \le n$. In this case, $r^n_{\mathcal{A}_x} = x^{\delta+i} + O(1)$.
- If $\delta \in (\alpha + \log_x 2, 1)$, then $x^{\delta+i} > n$. In this case, $r^n_{\mathcal{A}_x} = n - x^{\delta+i-1} + O(1)$ (note that $i \ge 1$ here).

Plugging the values above in the formula for the expected value, we obtain a different sum of integrals, which however leads to the same function as above:

$$\begin{aligned}\mathbb{E}[R_{\mathcal{A}_x}] &= \int_0^{\alpha} \frac{2x^{\alpha+i} - x^{\delta+i}}{x^{i+\alpha}} d\delta + \int_{\alpha}^{\alpha+\log_x 2} \frac{x^{\delta+i}}{x^{i+\alpha}} d\delta + \int_{\alpha+\log_x 2}^{1} \frac{2x^{\alpha+i} - x^{\delta+i-1}}{x^{i+\alpha}} d\delta + O(n^{-1}) \\ &= 2 - 2\log_x 2 + \frac{2}{x \ln x} + O(n^{-1}).\end{aligned}$$

Moreover, the formulas above are independent of α. Thus, it remains to find a value of x that minimizes $f(x) \stackrel{def}{=} 2 - 2\log_x 2 + \frac{2}{x \ln x}$. Observe that $f'(x) = -\frac{2}{x^2 \ln^2 x} - \frac{2}{x^2 \ln x} + \frac{\ln 2}{x \ln^2 x}$, and $f'(x) = 0$ if and only if $x = \log_2(ex)$. With a simple transformation, the latter is equivalent to $ze^z = -\frac{\ln 2}{e}$ with $z = -x \ln 2$, so the value that minimizes $f(x)$ can be computed as $x_{\min} = -\frac{W_{-1}(-\ln 2/e)}{\ln 2} \approx 3.052$, where W_{-1} is the lower branch of Lambert's W function. The claim of the theorem follows by calculating $f(x_{\min})$. □

Next, we prove that no algorithm can achieve an (expected) competitive ratio better than the one claimed in Theorem 2. The proof applies Yao's Minimax principle and uses a hard input distribution over all-ones sequences of length $n \in [n_{\min}, n_{\max}]$, for some large values of $n_{\min}$ and $n_{\max}$, where the probability that the sequence is of length n is proportional to $1/n$.

Theorem 3. For any randomized algorithm $\mathcal{A}$, $\mathbb{E}[R^n_{\mathcal{A}}] \geq 1.344$.

Proof. We will prove the theorem by using Yao's Minimax principle [26]. To this end, let us first consider an arbitrary *deterministic* algorithm $\mathcal{A}_{\text{det}}$. Assume the length of the sequence is random in the interval $X := [n_{\min}, n_{\max}]$ for large values of $n_{\max}$ and $n_{\min}$ with $n_{\max} > 2 \cdot n_{\min}$ and has the following distribution: The sequence ends at position $n \in X$ with probability p_n which is proportional to $\frac{1}{n}$, i.e., using the definition $S = \sum_{m=n_{\min}}^{n_{\max}} \frac{1}{m}$, we have

$$p_n := \mathbb{P}[\text{sequence is of length } n] = \frac{1}{n \cdot S} .$$

We will show that for *each* deterministic algorithm $\mathcal{A}_{\text{det}}$, if the input sequence is distributed as above, then $\mathbb{E}[R^n_{\mathcal{A}_{\text{det}}}] \geq 1.344 - O(\ln^{-1} n_{\max}/n_{\min})$, where the expectation is taken over the distribution of n.

Let J denote the set of requests at which $\mathcal{A}_{\text{det}}$ places the separator when processing the all-ones sequence of length $n_{\max}$. Note that the set of separator positions placed by $\mathcal{A}_{\text{det}}$ on sequences of shorter lengths are a subset of J. Let $I = J \cap X = \{i_1, \ldots, i_k\}$ (the i_j are ordered with increasing value).

For $n \in X$, let $r^n_{\mathcal{A}_{det}}$ be the bottleneck value of the partitioning computed by $\mathcal{A}_{\text{det}}$ on the sequence of length n. We bound $\mathbb{E}[R^n_{\mathcal{A}_{\text{det}}}] = \sum_{n=n_{\min}}^{n_{\max}} p_n R^n_{\mathcal{A}_{\text{det}}}$ by separately bounding every partial sum in the following decomposition:

$$\mathbb{E}[R^n_{\mathcal{A}_{\text{det}}}] = E(n_{\min}, i_1) + E(i_1, i_2) + \cdots + E(i_{k-1}, i_k) + E(i_k, n_{\max}),$$

where for each $a > b$, $E(a, b) = \sum_{n=a}^{b-1} p_n R^n_{\mathcal{A}_{\text{det}}}$. The first and last terms need a special care, so we will start with bounding all other terms. In the following, $H^q_p = \sum_{n=p}^{q} \frac{1}{n}$ denotes partial harmonic sums for $q \geq p \geq 1$. In particular, $S = H^{n_{\max}}_{n_{\min}}$.

Thus, we proceed in three steps:

1. Consider an index $1 \leq j < k$ and let us bound the sum $E(i_j, i_{j+1})$. Let us denote $a = i_j$ and $b = i_{j+1}$. We need to consider two cases.

   Case 1: $b \leq 2a$. Then for all $n \in \{a, \ldots, b-1\}$, the bottleneck value computed by the algorithm is $r^n_{\mathcal{A}_{det}} = a$ (since $n/2 < a$). Then:

   $$E(a, b) \geq \sum_{n=a}^{b-1} \frac{1}{nS} \cdot \frac{a}{\lceil n/2 \rceil} \geq \frac{2a}{S} \sum_{n=a}^{b-1} \frac{1}{n(n+1)} = \frac{2a}{S} \left(\frac{1}{a} - \frac{1}{b} \right) > 1.4 \cdot \frac{H^{b-1}_a}{S}, \tag{1}$$

   where the last inequality can be proved as follows. First, it is easily checked that the ratio $\Phi(a, b) = 2(1 - \frac{a}{b})/H^{b-1}_a$ decreases when b increases with a kept fixed, implying that (recall that $b \leq 2a$ and using standard approximations of harmonic sums)

$$\Phi(a,b) \geq \Phi(a,2a) = 1/H_a^{2a-1} \geq \frac{1}{\ln 2 + \frac{1}{a}} > 1.4,$$

where the last inequality holds when $a = i_j$ is large enough (say $i_j \geq n_{\min}/2 \geq 50$).

Case 2: $b > 2a$. In this case, for all $n = a, \ldots, 2a-1$, if the sequence is of length n, then $r^n_{\mathcal{A}_{det}} = a$, as in case 1. However, when $n \geq 2a$, then $n/2 \geq a$, so $r^n_{\mathcal{A}_{\text{det}}} = n - a$. Using these observations, we can bound $\mathbb{E}[R^n_{\mathcal{A}_{\text{det}}}]$ as follows:

$$\begin{aligned}
\mathbb{E}[R^n_{\mathcal{A}_{\text{det}}}] &= \sum_{n=a}^{2a-1} \frac{1}{nS}\frac{a}{\lceil n/2 \rceil} + \sum_{n=2a}^{b-1} \frac{1}{nS}\frac{n-a}{\lceil n/2 \rceil} \\
&\geq \frac{2a}{S}\sum_{n=a}^{2a-1} \frac{1}{n(n+1)} + \frac{2}{S}H^b_{2a+1} - \frac{2a}{S}\sum_{n=2a}^{b-1}\frac{1}{n(n+1)} \\
&= \frac{2a}{S}\left(\frac{1}{a} - \frac{1}{2a}\right) + \frac{2}{S}H^b_{2a+1} - \frac{2a}{S}\left(\frac{1}{2a} - \frac{1}{b}\right) \\
&= \frac{2}{S} + \frac{2a}{Sb} + \frac{2}{S}(H^b_a - H^{2a}_a) \\
&= \frac{H^b_a}{S}\cdot\left(2 + \frac{2a}{bH^b_a} - \frac{H^{2a}_a}{H^b_a}\right),
\end{aligned}$$

where the third line is obtained by using the identity $H^b_{2a+1} = H^b_a - H^{2a}_a$. Again, using a standard approximation for the harmonic sums, and setting $x = \frac{b}{a}$, we can approximate:

$$2 + \frac{2a}{bH^b_a} - \frac{H^{2a}_a}{H^b_a} \geq 2 + \frac{2a}{b\ln\frac{b}{a}} - \frac{\ln 2}{\ln\frac{b}{a}} - \mathrm{O}(a^{-1}) = 2 + \frac{2}{x\ln x} - \log_x 2 - \mathrm{O}(a^{-1}),$$

Note that the function $f(x) = 2 + \frac{2}{x\ln x} - \log_x 2$ is exactly the same that was minimized in the proof of Thm. 2, and achieves its minimum in $(1,\infty)$ at $x_{\min} \approx 3.052$, giving $f(x_{\min}) \approx 1.344$. Thus, we have $E(i_j, i_{j+1}) \geq \frac{H^{i_{j+1}-1}_{i_j}}{S}(1.344 - O(\frac{1}{i_j}))$.

2. The term $E(i_k, n_{\max})$ can be bounded by $\frac{H^{n_{\max}-1}_{i_k}}{S}(1.344 - O(\frac{1}{i_k}))$ by an identical argument as above.

3. The term $E(n_{\min}, i_1)$ needs a slightly different approach. Let i_0 denote the last separator that the algorithm placed before $n_{\min}$. We can assume that $i_0 \geq n_{\min}/2$, as otherwise the algorithm could only profit by moving the separator to $n_{\min}/2$. We consider two cases. First, if $i_1 \leq 2i_0$, then we simply assume the algorithm performs optimally in the range $[n_{\min}, i_1)$:

$$E(n_{\min}, i_1) = \frac{H^{i_1-1}_{n_{\min}}}{S} \geq 1.344\cdot\frac{H^{i_1-1}_{n_{\min}}}{S} - 0.5\frac{H^{i_1-1}_{n_{\min}}}{S} > 1.344\cdot\frac{H^{i_1-1}_{n_{\min}}}{S} - 1/S,$$

since (recalling that $i_1 \leq 2i_0 \leq 2n_{\min}$), $H^{i_1}_{n_{\min}} < H^{2n_{\min}}_{n_{\min}} < 1$.

On the other hand, when $i_1 > 2i_0$ (and by the discussion above, $2i_0 \geq n_{\min}$), then with calculations similar to the one in Case 2 above, we can obtain:

$$\begin{aligned}
E(x_{\min}, i_1) &= \sum_{n=n_{\min}}^{2i_0-1}\frac{1}{nS}\frac{i_0}{\lceil n/2\rceil} + \sum_{n=2i_0}^{i_1-1}\frac{1}{nS}\frac{n-i_0}{\lceil n/2\rceil} \\
&\geq \frac{2i_0}{S}\sum_{n=n_{\min}}^{2i_0-1}\frac{1}{n(n+1)} + \frac{2}{S}H^{i_1}_{2i_0+1} - \frac{2i_0}{S}\sum_{n=2i_0}^{i_1-1}\frac{1}{n(n+1)} \\
&\geq \frac{2}{S}\left(\frac{i_0}{n_{\min}} + \frac{i_0}{i_1} + H^{i_1}_{n_{\min}} - H^{2i_0}_{n_{\min}}\right) \geq 2\cdot\frac{H^{i_1}_{n_{\min}}}{S} - O(1/S),
\end{aligned}$$

because $i_0 \leq n_{\min}$ and thus $H^{2i_0}_{n_{\min}} < 1$.

It remains to plug the obtained estimates in (1):

$$\begin{aligned}\mathbb{E}[R^n_{\mathcal{A}_{\text{det}}}] &= E(n_{\min}, i_1) + \sum_{j=1}^{k-1} E(i_j, i_{j+1}) + E(i_k, n_{\max}) \\ &\geq (1.344 - O(1/n_{\min})) \cdot \frac{H^{i_1}_{n_{\min}} + \sum_{j=1}^{k-1} H^{i_{j+1}-1}_{i_j} + H^{n_{\max}}_{i_k}}{S} - O(1/S) \\ &= 1.344 - O(1/S).\end{aligned}$$

Last, by Yao's principle, every randomized algorithm has a competitive ratio of at least $\mathbb{E}[R^n_{\mathcal{A}_{\text{det}}}]$. □

4.2 General Weights

Algorithm $\mathcal{A}_x$ can be adapted to weighted sequences $X = w_1, w_2, \ldots, w_n$ of positive integers as follows: X can be thought of as a sequence of $W_n = \sum_{i=1}^n w_i$ unit weight requests and we simulate $\mathcal{A}_x$ on this unit weight sequence. Whenever $\mathcal{A}_x$ attempts to place a separator, but the position does not fall at the end of a weight w_i, the separator is placed after w_i.

If the weights of the sequence are bounded, algorithm $\mathcal{A}_x$ can be analyzed similarly as Theorem 2, by treating all requests as unit weights. This introduces an additional error term:

Corollary 1. There is a constant value of $x \approx 3.052$ such that $E[R^n_{\mathcal{A}_x}] \approx 1.344 + O(B/W_n)$ for any sequence $X = w_1, w_2, \ldots, w_n$ with weights $w_i \leq B$.

When arbitrary weights are possible, a non-trivial bound can still be proved. Interestingly, the optimal gap size between the separator positions is larger than in the case of unweighted sequences.

Theorem 4. There is a constant value of $x \approx 5.357$ such that $\mathbb{E}[R^X_{\mathcal{A}_x}] \leq 1.627 + O(W_n^{-1})$ for each sequence of total weight W_n.

Proof. Let $X = w_1, w_2, \ldots, w_n$ be the input sequence of total weight W_n, and let $m = \arg\min_{m'} \sum_{i \leq m'} w_i \geq \frac{W_n}{2}$. Then, w_m is the *central weight* of the sequence, and we denote $\frac{W_n}{2}$ the *central point*. We will argue first that replacing all w_i left of w_m by a sequence of $\sum_{i<m} w_i$ unit requests, and replacing all w_i right of w_m by a single large request of weight $\sum_{i>m} w_i$ worsens the approximation factor of the algorithm. Indeed, suppose that the algorithm attempts to place a separator at a position j that falls on an element w_i, which is located left of w_m. Then the algorithm places the separator after w_i, which brings the separator closer to the center and thus improves the partitioning. Similarly, suppose that the algorithm attempts to place a separator at position j that falls on an element w_i, which is located right of w_m. By replacing all weights located to the right of w_m by a single heavy element, the algorithm has to place the separator at the end of the sequence, which gives the worst possible approximation ratio. Thus, we suppose from now on that X is of the form $X = 1 \ldots 1 w_m B$, where B may be non-existent.

Assume that the central point splits the request with weight w_m into two parts $w', w'' \geq 0$, such that $w' + \sum_{i<m} w_i = \frac{W_n}{2}$. Clearly, the optimal bottleneck value is $opt = \frac{W_n}{2} + \min\{w', w''\}$. Further assume that $W_n = 2x^{i+\alpha}$ for $\alpha \in [0, 1)$ and $i \in \mathbb{N}$. Then, $\frac{W_n}{2} = x^{i+\alpha}$ is the central point. Let α_1, α_2 be such that $x^{i+\alpha_1}$ is the starting point of weight w_m, and $x^{i+\alpha_2}$ is the starting point of weight B. Note that α_2 is non-negative, but α_1 can be negative. However, if $\alpha_1 < 0$, we can replace $\alpha_1 = 0$ without decreasing the approximation ratio, because in both cases the algorithm places the separator after w_m, while in the case of negative α_1 the optimum can only be larger than when $\alpha_1 = 0$. Thus, we assume w.l.o.g. that $\alpha_1 \geq 0$, so $\alpha_1 \in [0, \alpha]$ and $\alpha_2 \in [\alpha, \alpha + \log_x 2]$.

Again, we consider several cases. In the estimates below, we ignore the rounding terms as they are all $O(1)$, and hence the error term in the approximation ratio is $O(opt^{-1}) = O(W_n^{-1})$. Also, we use W in place of W_n for brevity.

Case 1: $\alpha < 1 - \log_x 2$. In this case, $x^{i+1} > W$. So for all $\delta \in [0, \alpha + \log_x 2]$, $x^{i+\delta}$ is the position where the last separator would be (in the unit weights case). However, when $\delta > \alpha + \log_x 2$ then $x^{i+\delta} > W$ and $x^{i+\delta-1} < W/2$ is the last separator.

Here we have:

- $\delta \in [0, \alpha_1]$: $r_{\mathcal{A}_x} = W - x^{i+\delta}$ (all unit weights before w_m).

- $\delta \in [\alpha_1, \alpha_2]$: the separator is placed after w_m, and $r_{\mathcal{A}_x} = W/2 + w''$.
- $\delta \in [\alpha_2, \alpha + \log_x 2]$: assume the worst case bound $r_{\mathcal{A}_x} \leq W$.
- $\delta \in [\alpha + \log_x 2, 1]$: $r_{\mathcal{A}_x} = W - x^{i+\delta-1}$, as the last separator is at $x^{i+\delta-1} < x^i \leq x^{i+\alpha_1}$.

Computing the expectation gives:

$$\begin{aligned} E[R^X_{\mathcal{A}_x}] &\leq \int_0^{\alpha_1} \frac{W - x^{i+\delta}}{opt} d\delta + \int_{\alpha_1}^{\alpha_2} \frac{W/2 + w''}{opt} d\delta + \int_{\alpha_2}^{\alpha+\log_x 2} \frac{W}{opt} d\delta \\ &+ \int_{\alpha+\log_x 2}^{1} \frac{W - x^{i+\delta-1}}{opt} d\delta \\ &= \frac{1}{opt} \left(W(1 + \frac{1}{x \ln x}) - \alpha_2(W/2 - w'') + \alpha_1(W/2 - w'') - \frac{x^{i+\alpha_1}}{\ln x} \right). \end{aligned} \tag{2}$$

First, assume that $w'' \leq w'$, which implies that $\alpha_1 \leq \log_x(W/2 - w'') - i$. In this case $opt = W/2 + w''$. Let us see which values of α_1 maximize the term $\phi(\alpha_1) = (\frac{W}{2} - w'')\alpha_1 - \frac{x^{i+\alpha_1}}{\ln x}$ in the parentheses. We have $\phi'(\alpha_1) = \frac{W}{2} - w'' - x^{i+\alpha_1} \geq 0$, since $\alpha_1 \leq \log_x(W/2 - w'') - i$ by assumption (that $w'' \leq w'$). So $\phi(\alpha_1)$ is increasing in the interval $[0, \log_x(W/2 - w'') - i]$, and $\alpha_1 = \log_x(W/2 - w'') - i$ is the maximizer. Plugging this value and $\alpha_2 = \log_x(W/2 + w'') - i$ in the expectation formula and rearranging the terms gives:

$$E[R^X_{\mathcal{A}_x}] \leq \frac{W(1 + \frac{1}{x \ln x}) + (\frac{W}{2} - w'')(\log_x(\frac{W}{2} - w'') - \log_x(\frac{W}{2} + w'') - \frac{1}{\ln x})}{\frac{W}{2} + w''}.$$

Note that the right hand side is a decreasing function of w'' and the maximum is achieved when $w'' = 0$:

$$E[R^X_{\mathcal{A}_x}] \leq 2 + \frac{2}{x \ln x} - \frac{1}{\ln x}. \tag{3}$$

Now assume that $w'' \geq w'$. Then $opt = W/2 + w'$. By plugging the value $\alpha_2 = \log_x(W/2 + w'') - i$ in (2) we see again that the expression is a decreasing function of w'', so it is maximum when w'' takes its minimum value $w'' = w'$:

$$E[R^X_{\mathcal{A}_x}] \leq \frac{W(1 + \frac{1}{x \ln x}) + (\frac{W}{2} - w')(\log_x(\frac{W}{2} - w') - \log_x(\frac{W}{2} + w') - \frac{1}{\ln x})}{\frac{W}{2} + w'}.$$

where we also used $x^{\alpha_1 + i} = W/2 - w'$. This is again a decreasing function of w' and gives exactly the same bound (3).

Case 2: $\alpha > 1 - \log_x 2$. Here $x^{i+\delta+1} < W$ for all $\delta \in [0, \alpha + \log_x 2 - 1]$, so we assume that the algorithm places the separator at the end of the sequence for such δ. Here we have to distinguish two sub-cases.

Case 2.1: $\alpha_1 > \alpha + \log_x 2 - 1$. We have:

- $\delta \in [0, \alpha + \log_x 2 - 1]$: $r_{\mathcal{A}_x} \leq W$.
- $\delta \in [\alpha + \log_x 2 - 1, \alpha_1]$: $r_{\mathcal{A}_x} = W - x^{i+\delta}$.
- $\delta \in [\alpha_1, \alpha_2]$: $r_{\mathcal{A}_x} = W/2 + w''$.
- $\delta \in [\alpha_2, 1]$: $r_{\mathcal{A}_x} \leq W$.

Computing the expectation gives:

$$\begin{aligned} E[R^X_{\mathcal{A}_x}] &\leq \int_0^{\alpha+\log_x 2-1} \frac{W}{opt} d\delta + \int_{\alpha+\log_x 2-1}^{\alpha_1} \frac{W - x^{i+\delta}}{opt} d\delta + \int_{\alpha_1}^{\alpha_2} \frac{W/2 + w''}{opt} d\delta + \int_{\alpha_2}^{1} \frac{W}{opt} d\delta \\ &= \frac{1}{opt} \left(W(1 + \frac{1}{x \ln x}) - \alpha_2(W/2 - w'') + \alpha_1(W/2 - w'') - \frac{x^{i+\alpha_1}}{\ln x} \right). \end{aligned}$$

The latter is the same formula as in Case 1, so (3) holds in this case too.

Case 2.2: $\alpha_1 \leq \alpha + \log_x 2 - 1$. In the analysis below, we assume that $r_{\mathcal{A}_x}$ takes the worst-case value W for all $\delta \in [0, \alpha + \log_x 2 - 1]$. In this case, we assume w.l.o.g. that $\alpha_1 = \alpha + \log_x 2 - 1$, because the optimum can only become worse when α_1 is smaller, while the algorithm (as we assume) will not profit. Thus, $x^{i+\alpha_1} = W/x$ and $\alpha_1 = \log_x W - 1 - i$. We have:

- $\delta \in [0, \alpha_1]$: $r_{\mathcal{A}_x} \leq W$.
- $\delta \in [\alpha_1, \alpha_2]$: $r_{\mathcal{A}_x} = W/2 + w''$.
- $\delta \in [\alpha_2, 1]$: $r_{\mathcal{A}_x} \leq W$.

Computing the expectation gives:

$$\begin{aligned} E[R^X_{\mathcal{A}_x}] &\leq \int_0^{\alpha_1} \frac{W}{opt} d\delta + \int_{\alpha_1}^{\alpha_2} \frac{W/2 + w''}{opt} d\delta + \int_{\alpha_2}^{1} \frac{W}{opt} d\delta \\ &= \frac{1}{opt} \left(W + \alpha_1 (W/2 - w'') - \alpha_2 (W/2 - w'') \right) \\ &= \frac{1}{opt} \left(W + (W/2 - w'') (\log_x \frac{W}{W/2 + w''} - 1) \right), \end{aligned} \tag{4}$$

where we used $\alpha_1 = \log_x W - 1 - i$ and $\alpha_2 = \log_x (W/2 + w'') - i$.

Now, if $w'' \leq w'$, then $opt = W/2 + w''$. Plugging this value in (4), we obtain the expression $(Z-1)\log_x Z + 1$ with $Z = \frac{W}{W/2+w''}$, which is a decreasing function of w'', so it is maximized for $w'' = 0$, in which case we have $E[R^X_{\mathcal{A}_x}] \leq 1 + \log_x 2$.

On the other hand, if $w'' > w'$, then $opt = W/2 + w' = W/2 + W/2 - x^{i+\alpha_1} = W - W/x$. Let us optimize the numerator of (4) as a function of w''. The derivative is

$$-\log_x \frac{W}{W/2 + w''} + 1 + \frac{W/2 - w''}{(W/2 + w'') \ln x} = 1 - \frac{1}{\ln x} \left(\frac{W}{W/2 + w''} + \ln \frac{W}{W/2 + w''} \right)$$

Recall that $w'' > w' = W/2 - W/x$, so the expression above is at least $1 - \frac{1}{\ln x} \left(\frac{1}{1-1/x} + \ln \frac{1}{1-1/x} \right)$ and is positive for all $x > 3$, showing that the bound we obtained is an increasing function of w'' and is maximized when $w'' = W/2$. Plugging the values of opt and $w'' = W/2$ in (4), we obtain that $E[R^X_{\mathcal{A}_x}] \leq \frac{1}{1/2+1/x}$.

Thus, covering all cases, we obtain the bound $E[R^X_{\mathcal{A}_x}] \leq \max\{g(x), h(x), f(x)\} + O(W_n^{-1})$, where $g(x) \stackrel{def}{=} 2 + \frac{2}{x \ln x} - \frac{1}{\ln x}$, $h(x) = 1 + \log_x 2$, and $f(x) = \frac{1}{1/2+1/x}$. It can be shown as in Thm. 2 that $x_{min} = -2W_{-1}(-\frac{1}{2e}) \approx 5.3567$ minimizes $g(x)$ and $g(x_{min}) \approx 1.627$. It can also be checked that $\max\{h(x_{min}), f(x_{min})\} < 1.5$. □

The competitive ratio 1.627 is tight for algorithm $\mathcal{A}_x$. This is achieved on sequences consisting of $W_n/2$ unit weight requests followed by a request of weight $W_n/2$. Furthermore, using a similar idea as for the unweighted case, a barely-random algorithm that uses only a single bit can be obtained with competitive ratio 1.75 (analysis omitted).

Last, it can be seen that on sequences with exponentially increasing weights, no algorithm can achieve a competitive ratio better than 1.5.

Theorem 5. Every randomized algorithm for PART on sequences with arbitrary weights has an expected competitive ratio of at least 1.5.

Proof. Let $S_i = 2^0, 2^1, 2^2, \ldots, 2^{i-1}$ denote the exponentially increasing sequence of length i, let $x_{\min}, x_{\max}$ be large integers with $x_{\max} \geq 2x_{\min}$, and let $X = [x_{\min}, x_{\max}]$. We consider the performance of any deterministic algorithm $\mathcal{A}_{\text{det}}$ on the uniform input distribution on set $\mathcal{S} = \{S_i : i \in X\}$. The result for randomized algorithms then follows by applying Yao's lemma.

Let J be the set of requests at which $\mathcal{A}_{\text{det}}$ places a separator on input $S_{x_{\max}}$. Note that the set of separator positions on any other input of $\mathcal{S}$ is a subset of J. Let $I = J \cap X = \{i_1, \ldots, i_k\}$ be the separators placed within the interval X (ordered such that $i_j < i_{j+1}$, for every j) and let i_0 be the right-most separator placed before $x_{\min}$. W.l.o.g., we can savely assume that $i_0 = x_{\min} - 1$, since this

does not worsen the algorithm (it is the optimal choice for $S_{x_{\min}}$). Furthermore, we can also assume that $x_k < x_{\max}$, since placing a separator at request $x_{\max}$ gives the worst ratio possible for $S_{x_{\max}}$.

We bound now the expected competitive ratio of $\mathcal{A}_{\text{det}}$, where the expectation is taken over the inputs $\mathcal{S}$. To this end, notice that the optimal bottleneck value OPT_i on sequence S_i is $OPT_i = 2^{i-1}$. We write $R^i_{\mathcal{A}_{\text{det}}}$ to denote the competitive ratio of $\mathcal{A}_{\text{det}}$ on sequence S_i. Then:

$$\begin{aligned}
\mathbb{E}_{S_i \leftarrow \mathcal{S}} R^i_{\mathcal{A}_{\text{det}}} &= \frac{1}{x_{\max} - x_{\min} + 1} \cdot \sum_{n=x_{\min}}^{x_{\max}} R^n_{\mathcal{A}_{\text{det}}} \text{ , and} \\
\sum_{n=x_{\min}}^{x_{\max}} R^n_{\mathcal{A}_{\text{det}}} &= \underbrace{\sum_{n=x_{\min}}^{i_1-1} R^n_{\mathcal{A}_{\text{det}}}}_{I} + \sum_{n=i_1}^{i_2-1} R^n_{\mathcal{A}_{\text{det}}} + \cdots + \sum_{n=i_{k-1}}^{i_k-1} R^n_{\mathcal{A}_{\text{det}}} + \underbrace{\sum_{n=i_k}^{x_{\max}} R^n_{\mathcal{A}_{\text{det}}}}_{II} .
\end{aligned}$$

We now bound I, II, and $\sum_{n=i_j}^{i_{j+1}-1} R^n_{\mathcal{A}_{\text{det}}}$ for every $1 \le j \le k-1$, separately.

1. First, observe that for every $1 \le j \le k-1$, we have $R^{i_j}_{\mathcal{A}_{\text{det}}} = \frac{2^{i_j}-1}{2^{i_j-1}} = 2 - \frac{1}{2^{i_j-1}}$, $R^{i_j+1}_{\mathcal{A}_{\text{det}}} = 1$ (if $i_j + 1 < i_{j+1}$), and in general for every $a \ge 2$ with $i_j + a < i_{j+1}$, $R^{i_j+a}_{\mathcal{A}_{\text{det}}} = \frac{2^{i_j}+2^{i_j+1}+\cdots+2^{i_j+a-1}}{2^{i_j+a-1}} \ge 1.5$. Hence, $\sum_{n=i_j}^{i_{j+1}-1} R^n_{\mathcal{A}_{\text{det}}} \ge (1.5 - \frac{1}{2^{i_j-1}})(i_{j+1} - i_j) \ge (1.5 - \frac{1}{2^{x_{\min}-1}})(i_{j+1} - i_j)$.
2. Concerning I, first notice that if $i_1 = x_{\min}$, then $I = 0$. By similar considerations as above, if $i_1 = x_{\min} + a$, then $I = \sum_{b=0}^{a-1} \frac{2^{x_{\min}-1}+2^{x_{\min}+1}+\cdots+2^{x_{\min}+b-1}}{2^{x_{\min}+b-1}}$ which is at least $1.5 \cdot a$, if $a \ge 4$.
3. Last, concerning II, this case is identical to the first case, and we can bound II by $II \ge (1.5 - \frac{1}{2^{x_{\min}-1}})(x_{\max} - i_k)$.

Thus, we can bound $\mathbb{E}_{i \leftarrow X} R^i_{\mathcal{A}_{\text{det}}}$ by:

$$\begin{aligned}
\mathbb{E}_{S_i \leftarrow \mathcal{S}} R^i_{\mathcal{A}_{\text{det}}} &\ge \frac{1}{x_{\max} - x_{\min} + 1} \left((x_{\max} - x_{\min} + 1 - 4)(1.5 - \frac{1}{2^{x_{\min}-1}}) + 4 \cdot 1 \right) \\
&= 1.5 - \mathrm{O}(\frac{1}{x_{\max}}).
\end{aligned}$$

Thus, since $x_{\max}$ can be chosen arbitrarily large, every deterministic algorithm has an expected competitive ratio of 1.5. The result for randomized algorithms follows by applying Yao's principle. □

5 Partitioning with Arbitrary Number of Partitions

We give now our results for Part for arbitrary p. We first give an algorithm and a lower bound that directly consider Part. Then, we address Part indirectly, by first solving Flow.

5.1 Algorithm and Lower Bound for Part: Direct Approach

5.1.1 Algorithm

We give now a deterministic 2-competitive algorithm for Part for any number of partitions p in the preemptive online model.

A building block of our algorithm is the Probe algorithm, which has previously been used for tackling Part [7, 14, 16]. Algorithm Probe takes an integer parameter B, which constitutes a potential bottleneck value, and traverses the input sequence from left to right, placing separators such that partitions of maximal size not larger than B are created. It is easy to see that Probe creates at most p partitions if $B \ge B^*$, the optimal bottleneck value.

Lemma 1. Let X be an integer sequence. If there exists a partitioning of X into p parts with maximum partition weight m, then Probe(m) creates at most p partitions.

Algorithm 3 2-Approximation for PART

$w_i \leftarrow 0$ for all $1 \leq i \leq p$ {current partitions}
$S \leftarrow 0$ {current total weight}
$m \leftarrow 0$ {current maximum}
while request sequence not empty **do**
  $x \leftarrow$ next request
  $S \leftarrow S + x$, $m \leftarrow \max\{m, x\}$
  $B \leftarrow 2 \cdot \max\{m, S/p\}$ {bottleneck val.}
  Run PROBE(B) on $w_1, w_2, \ldots, w_p, x$ and update partition weights $w_1, \ldots, w_p$
end while

Instead of running PROBE directly on the input sequence, we will run PROBE for a carefully chosen bottleneck value B on the sequence $w_1, \ldots, w_p, x$ of current partition weights $w_1, \ldots, w_p$ followed by the value of the current request x. We will prove that this run of PROBE creates at most p partitions. If PROBE places a subsequence $w_i, \ldots, w_j$ of partition weights into the same partition, then the partition separators between current partitions $i, \ldots, j$ are removed. See Algorithm 3 for details.

Theorem 6. Algorithm 3 is a deterministic 2-competitive algorithm for PART.

Proof. First, suppose that the run of PROBE Algorithm 3 succeeds in every iteration. Let S be the weight of the entire input sequence, and let m be its maximum. Since the optimal bottleneck value B^* is trivially bounded from below by $\max\{m, S/p\}$, and the bottleneck value employed in the last run of PROBE is $B = 2 \cdot \max\{m, S/p\}$, we obtain an approximation factor of 2.

Denote $w_{p+1} = x$. It remains to prove that the run of PROBE always succeeds, i.e., in every iteration of the algorithm, the optimal bottleneck value of the sequence $w_1, w_2, \ldots w_p, w_{p+1}$ is at most $B = 2 \cdot \max\{m, S/p\}$ (where S is the current total weight of the input sequence, and m the current maximum). Indeed, if PROBE(B) does not succeed in creating p partitions, then $w_i + w_{i+1} > B \geq 2S/p$ must hold for all $1 \leq i \leq p$. But then:

$$S = \sum_{i=1}^{p+1} w_i \geq \sum_{i=1}^{\lfloor (p+1)/2 \rfloor} (w_{2i-1} + w_{2i}) > \lfloor (p+1)/2 \rfloor \cdot 2S/p \geq S,$$

a contradiction, which proves correctness of the algorithm. □

5.1.2 Lower Bound

Next, we present a lower bound for PART that makes use of the discrete properties of integers (and does not apply to FLOW). For a given $p \geq 2$, let σ_1 and σ_2 be the all-ones sequences of lengths $2p$ and $2p+1$, respectively. Consider first a deterministic algorithm $\mathcal{A}_{\text{det}}$. If $\mathcal{A}_{\text{det}}$ computes an optimal solution on σ_1, i.e., a partitioning consisting of p partitions each of weight 2, then $\mathcal{A}_{\text{det}}$ on σ_2 is at least $\frac{4}{3}$-competitive, since the merging of any two partitions creates a bottleneck value of 4 while the optimal bottleneck is 3. On the other hand, if $\mathcal{A}_{\text{det}}$ is not optimal on σ_1, then $\mathcal{A}_{\text{det}}$ is at least $\frac{3}{2}$-competitive on σ_1. Thus, every deterministic algorithm for PART is at least $\frac{4}{3}$-competitive.

Consider now a randomized algorithm $\mathcal{A}_{\text{rand}}$. Consider further the input distribution over $\{\sigma_1, \sigma_2\}$ so that σ_1 occurs with probability $\frac{2}{5}$, and σ_2 occurs with probability $\frac{3}{5}$. Let p denote the probability that $\mathcal{A}_{\text{rand}}$ outputs an optimal solution on σ_1. Then, the approximation ratio of $\mathcal{A}_{\text{rand}}$ is bounded from below by:

$$\mathbb{P}[\sigma_1 \text{ occurs}] \cdot \left(p \cdot 1 + (1-p) \cdot \frac{3}{2} \right) + [\sigma_2 \text{ occurs}] \cdot \left(p \cdot \frac{4}{3} + (1-p) \cdot 1 \right) = \frac{6}{5} .$$

We thus established the following theorem:

Theorem 7. Every deterministic (randomized) preemptive online algorithm for PART is at least $\frac{4}{3}$-competitive (resp. $\frac{6}{5}$-competitive).

Algorithm 4 Periodic Scheme

Require: Integer p, real number x
  $S_i \leftarrow$ super-partition with base i, $\forall i \in [p]$
  $j \leftarrow 1$
  $X_j \leftarrow S_1$ {Initial conf. equals S_1}
  for $p-1$ times **do**
    for $i \leftarrow 1 \dots p$ **do**
      merge-next(S_i)
      $X' \leftarrow$ length p prefix of $S_1, S_2, \dots, S_p$
      if $X' \neq X_j$ **then**
        $j \leftarrow j+1$, $X_j \leftarrow X'$
      end if
    end for
  end for

5.2 Algorithm and Lower Bound for Part: Indirect Approach Via Flow

5.2.1 Algorithm

We will first give a deterministic algorithm for Flow, which is roughly 1.68-competitive, and then prove a lower bound of 1.08.

In the following, we assume that p is a power of two. We first give a scheme, which is based on a positive real number x (we will set $x = 2$ later), that defines a sequence of p-ary vectors $X_1, X_2, \dots$. Then, we will prove in Lemma 2 how this scheme can be applied to Flow. To define our scheme, we require the concept of a *super-partition*:

Definition 3 (Super-partition)**.** For an integer $1 \leq b \leq p$ denoted the base, let $S = S_1, \dots, S_p$ with $S_1 = x^{\frac{b}{p}}$ and $S_{i+1} = x^{\frac{1}{p}} S_i$ be the initial configuration of the super-partition. The super-partition evolves by merging in each step the pair of partitions whose sum is minimal, thus decreasing the length of the super-partition by 1 in every step (since we assumed that p is a power of two, overall we do $\log n$ merge sweeps from left to right).

Our scheme is best explained via the evolution of p super-partitions with different bases and is depicted in Algorithm 4. It outputs a sequence $X_1, X_2, \dots$, which corresponds to the evolution of the weights of the p partitions in an algorithm for Flow.

To illustrate the scheme, as in Table 1 in the introduction, we consider the case $p = 4$. Initially, four super-partitions are in their initial state, and since X_1 equals the length 4 prefix of the values of the super-partitions $S_1, S_2, \dots, S_p$, we have that X_1 equals S_1 (in the following, the X_i are highlighted in bold):

$$\overbrace{\boldsymbol{x^{\frac{1}{4}}, x^{\frac{2}{4}}, x^{\frac{3}{4}}, x^{\frac{4}{4}}}}^{S_1} \quad | \quad \overbrace{x^{\frac{2}{4}}, x^{\frac{3}{4}}, x^{\frac{4}{4}}, x^{\frac{5}{4}}}^{S_2} \quad | \quad \overbrace{x^{\frac{3}{4}}, x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}}}^{S_3} \quad | \quad \overbrace{x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}}, x^{\frac{7}{4}}}^{S_4}$$

Then, merge-next advances super-partition S_1 into its next state, by merging the two lightest weights ($x^{\frac{1}{4}}$ and $x^{\frac{2}{4}}$ to $x^{\frac{1}{4}} + x^{\frac{2}{4}} = x^{\frac{1}{4}}(1 + x^{\frac{1}{4}})$). This makes the weight $x^{\frac{2}{4}}$ of super-partition S_2 advance to position 4 and is thus included in X_2:

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{3}{4}}, x^{\frac{4}{4}}} \quad | \quad \boldsymbol{x^{\frac{2}{4}}}, x^{\frac{3}{4}}, x^{\frac{4}{4}}, x^{\frac{5}{4}} \quad | \quad x^{\frac{3}{4}}, x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}} \quad | \quad x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}}, x^{\frac{7}{4}}$$

Next, super-partition S_2 is advanced, which gives X_3:

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{3}{4}}, x^{\frac{4}{4}}} \quad | \quad \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}})}, x^{\frac{4}{4}}, x^{\frac{5}{4}} \quad | \quad x^{\frac{3}{4}}, x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}} \quad | \quad x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}}, x^{\frac{7}{4}}$$

Note that in the next two steps, super-partitions S_3 and S_4 are advanced, which does not affect the first four positions. Hence, in these iterations, the condition in Line 8 evaluates to false and no new X_j values are created. The entire evolution is illustrated in Appendix B.

Analysis. Throughout the analysis, we use the notation $X_i[j]$ to denote the jth element of X_i. We also fix $x = 2$. Some of our results are stated for $x = 2$ while others use a general x for convenience. We also

use the abbreviation $\alpha = x^{1/p} = 2^{1/p}$. We first show how our scheme can be used to obtain an algorithm for FLOW.

Lemma 2. The scheme of Algorithm 4 gives rise to a deterministic algorithm for FLOW with competitive ratio $\frac{p+4}{p} \cdot \max_j \frac{\max X_j}{\operatorname{avg} X_j}$.

Proof. W.l.o.g., we assume that the warm-up period in FLOW is time $\sum_{x \in S_1} x$ (if not, then we scale our scheme). We initialize the p partitions with X_1. The transition from X_i to X_{i+1} in our scheme should be understood as a two-step process: First, the merge operation takes place (i.e., the advancing of some super-partition), which either creates an empty partition to the right, or the last partition has been merged with more incoming flow. In case an empty partition is created, it is then filled with flow up to value $X_{i+1}[p]$. When our scheme terminates but the flow has not yet ended, then we simply repeat the scheme, where all values are scaled. To see that this is possible, let $X_{\max}$ be the last partition created by Algorithm 4, and observe that $X_{\max}[i] = x^{\frac{i-1}{p}} \sum_{j=1}^{p} S_1[j]$, while $X_1[i] = x^{\frac{i}{p}}$. Hence, the final configuration is a scaled version of the initial configuration (by factor $x^{-\frac{1}{p}} \sum_{j=1}^{p} S_1[j]$).

It is therefore enough to assume that the flow stops during the first iteration of the scheme. At that moment, the last partition may not be entirely filled. Consider a final partitioning such that the first $p-1$ partitions coincide with the weights of the first $p-1$ values of some X_i, and the last partition is of some arbitrary weight $y \leq X_i[p]$. Then, the competitive ratio is bounded by:

$$\frac{\max\{X_i[1], \ldots, X_i[p-1], y\}}{\frac{1}{p}\left(y + \sum_{j=1}^{p-1} X_i[j]\right)} \leq \frac{\max X_i}{\operatorname{avg} X_{i-1}} \leq \frac{p+4}{p} \cdot \frac{\max X_i}{\operatorname{avg} X_i} \ ,$$

where we used $\operatorname{avg} X_{i-1} \geq \frac{p+4}{p} \operatorname{avg} X_i$, which will be proved below. □

The next lemma was used in the proof of the previous lemma.

Lemma 3. Let $x = 2$. Then, for every i:

$$\operatorname{avg} X_{i-1} \geq \frac{p}{p+4} \operatorname{avg} X_i.$$

Proof. First, by investigating the structure of X_i, it can be seen that $\max X_i \leq 2 \cdot \min X_i$, for every X_i, and $\max X_i \leq 2 \max X_{i-1}$. These two bounds give $\max X_i \leq 4 \min X_{i-1}$, which implies $\max X_i \leq 4 \operatorname{avg} X_{i-1}$. Then,

$$\begin{aligned} \operatorname{avg} X_i &= \frac{1}{p} \cdot \sum_{x \in X_i} x = \frac{1}{p} \cdot \left(X_i[p] + \sum_{y \in X_{i-1}} y \right) \leq \frac{1}{p} \cdot \left(4 \operatorname{avg} X_{i-1} + \sum_{y \in X_{i-1}} y \right) \\ &= \frac{p+4}{p} \operatorname{avg} X_{i-1}. \end{aligned}$$

□

Thus, in order to obtain a good algorithm for FLOW, we need to bound the max-over-average ratio $\frac{\max X_j}{\operatorname{avg} X_j}$, for every X_j, of the scheme. To this end, observe that X_i is of the form:

$$X_i = \underbrace{S_1 S_2 \ldots S_j}_{\text{Length } l}[\ \underbrace{S_{j+1} \ldots S_k}_{\text{Length } l+1}\ ],$$

where $S_1, \ldots, S_j$ are of length l, $S_{j+1} \ldots, S_k$ are of length $l+1$, $S_{j+1} \ldots, S_k$ may or may not exist, and the last super-partition of X_i may be incomplete (not entirely included in X_i). Using Lemma 4, it can be seen that for base $x = 2$, the maximum of X_i always lies in either the right-most super-partition of length l that is entirely included in S_i, or in the incomplete super-partition of length l (if it exists).

In general, in order to describe the current X_i, let L, t, m be such that L is a power of two, $1 \leq t \leq L/2$, all super-partitions have lengths either $l = L - t$ or $l + 1 = L - t + 1$, and m is the rightmost index of the super-partition that has length $L - t$ and is entirely included in X_i. Such a triple (L, t, m) completely describes the state of X_i.

First, we need an expression for $\max X_i$. This quantity depends on the current lengths of the super-partitions and the value m, and using the quantities L and l, we obtain the following lemma:

Our analysis requires a structural result on super-partitions, which we give first. In the following lemma, we relate the maximum value of a given super-partition to its length, and we give an expression for the total weight of a super-partition.

Lemma 4. Fix a super-partition S_i and suppose that its current length is $l = L - t$, where L is a power of two and $1 \le t \le L/2$. Let max_i denote the maximum at that moment and sum_i the total weight of S_i. Then:

$$sum_i = \frac{\alpha^i}{\alpha - 1}, \quad \text{and} \quad max_i = \alpha^i \cdot \alpha^{2tp/L} \cdot \frac{1 - \alpha^{-2p/L}}{\alpha - 1}.$$

Proof. The first claim is easy: We have

$$sum_1 = \sum_{j=1}^{p} \alpha^j = \frac{\alpha^{p+1} - \alpha}{\alpha - 1} = \frac{\alpha}{\alpha - 1}, \quad \text{and} \quad sum_i = \alpha^{i-1} \cdot sum_1 \ ,$$

which gives the result.

Denote by $s_{i,j}$ the jth element of S_i. Concerning the maximum, suppose first that super-partition S_1 is of length L. Then, the kth element of S_1 is $s_{1,k} = \alpha^{(k-1)p/L} \cdot s_{1,1} = \alpha^{(k-1)p/L} \cdot \sum_{j=1}^{p/L} \alpha^j = \alpha^{(k-1)p/L} \cdot \frac{\alpha^{p/L+1} - \alpha}{\alpha - 1}$. The corresponding elements in S_i, when it has length L, are $s_{i,k} = \alpha^{i-1} \cdot s_{1,k} = \alpha^{i-1} \cdot \alpha^{(k-1)p/L} \cdot \frac{\alpha^{p/L+1} - \alpha}{\alpha - 1}$. Consider now the general setting, when S_i is of length $l = L - t$. Then:

$$max_i = s_{i,2t-1} + s_{i,2t} = (1 + \alpha^{p/L}) \cdot \alpha^{(2t-2)p/L} \cdot \alpha^{i-1} \cdot \frac{\alpha^{p/L+1} - \alpha}{\alpha - 1} = \alpha^i \cdot \alpha^{2(t-1)p/L} \cdot \frac{\alpha^{2p/L} - 1}{\alpha - 1}.$$

□

Bounding $\operatorname{avg} X_i$ is the challenging part. Based on the triple (L, t, m), our analysis requires a bound on c, the number of super-partitions entirely included in X_i. It is easy to see that c is at least $\lceil \frac{p+m-l}{l+1} \rceil$. This, however, introduces difficulties, since optimizing over such a function is difficult if we need a fine enough optimization that does not allow us to ignore rounding effects (i.e., the estimate $c \ge \frac{p+m-l}{l+1}$ is not good enough in some cases).

In Lemma 5 we thus conduct a case distinction: In the easy cases (if c is large), the estimate $c \ge \frac{p+m-l}{l+1}$ is good enough for our purposes and the max-over-average ratio follows by a careful calculation. In the other cases (c is small), we analyze the max-over-average ratio using *reference points*: A reference point is one where all partitions included in X_i have equal length which is a power of two (for example the initial configuration is one). We then study the behavior of the scheme between two consecutive reference points and obtain a different bound. Lemma 5 is the most technical part of this analysis.

Lemma 5. Let $R = \frac{\max X_i}{\operatorname{avg} X_i}$. The following bounds hold for each X_i with configuration (L, t, m), with an additive error term in $O(1/p)$:

$$R \le \min \left\{ \frac{\ln 2 \cdot 2^{2t/L}(1 - 2^{-2/L})}{2^{\frac{p+1-l}{p(l+1)}} - 1}, \frac{2p \cdot (2^{2/L} - 1)L}{2^{1/\ln 2} \ln 2 \cdot (p + 2 - L)}, \frac{4 \ln 2 \cdot (2^{2/L} - 1) \cdot 2^{-\frac{2p}{cL}}}{2^{\frac{c+2}{p}}(2^{1/p} - 1) + 2^{-1/p} - 2^{-c/p}} \right\}.$$

Proof. The first and second bounds will be used when there are relatively many super-partitions in X_i. In that case, we can afford ignoring the remainder part S_{c+1} (when estimating the average) and also use the bound on c that ignores rounding effects.

Thus, the average is estimated as follows:

$$p \cdot \operatorname{avg} X_i \ge \sum_{t=1}^{c} sum_t = \frac{\sum_1^c \alpha^t}{\alpha - 1} = \frac{\alpha^{c+1} - \alpha}{(\alpha - 1)^2} \ .$$

Recall that the current maximum is achieved in S_m, so all we need is to bound the ratio $max_m / \operatorname{avg} X_i$. Using the expression for max_m given in Lemma 4, we obtain:

$$\frac{max_m}{\operatorname{avg} X_i} \le \frac{p(\alpha - 1)\alpha^{m-1}\alpha^{2tp/L}(1 - \alpha^{-2p/L})}{\alpha^c - 1} \le \frac{\ln 2 \cdot \alpha^{m-1}\alpha^{2tp/L}(1 - \alpha^{-2p/L})}{\alpha^{\frac{p+m}{l+1}-1} - 1},$$

where we used the approximation $p(1-\alpha) \approx \ln 2$ which holds with an error bounded by $1/p$ (which we will henceforth omit) and the bound $c \geq \frac{p+m}{l+1} - 1$. Differentiation shows that the right hand side is non-increasing as a function of m and is maximum at $m = 1$, which gives the first bound:

$$\frac{max_m}{\operatorname{avg} X_i} \leq \frac{\ln 2 \cdot \alpha^{2tp/L}(1 - \alpha^{-2p/L})}{\alpha^{\frac{p-l}{l+1}} - 1} .$$

This can be used to obtain bounds for smaller values of l. For larger values, we use the bound $e^x > 1 + x$, applied to $\alpha^{\frac{p-l}{l+1}} - 1$ in the denominator, which gives $\alpha^{\frac{p-l}{l+1}} - 1 = 2^{\frac{p-l}{p(l+1)}} - 1 > \frac{\ln 2 \cdot (p-l)}{p(l+1)}$, which gives, after replacing $\alpha = 2^{1/p}$,

$$\frac{max_m}{\operatorname{avg} X_i} \leq \frac{p \cdot (1 - 2^{-2/L}) \cdot 2^{2t/L}(L - t + 1)}{p - L + t} .$$

A simple and crude optimization w.r.t. t, where we replace $t = 1$ in the denominator and then optimize the numerator w.r.t. t (giving $t = \left(1 - \frac{1}{2 \ln 2}\right) L + 1$), gives the second bound:

$$\frac{max_m}{\operatorname{avg} X_i} \leq \frac{2p \cdot (2^{2/L} - 1)L}{2^{1/\ln 2} \ln 2 \cdot (p + 1 - L)} .$$

Now, consider the case when there are few super-partitions in X_i, i.e., c is small. We need a more delicate bound on $\operatorname{avg} X_i$ (by losing precision when computing $\max X_i$: uncertainty principle in action), so we shift the reference point to the time step when there are *exactly* c super-partitions in X_i and no incomplete super-partitions. Let (L, t_c, m_c) be the tuple describing this point. We concentrate on the process of going from (L, t_c, m_c) to (L, t_{c+1}, m_{c+1}). Note that $(L - t_c)m_c + (c - m_c)(L - t_c + 1) = p$, and $0 < m_c < c$, so we have $L - \frac{p}{c} + \frac{1}{c} \leq t_c \leq L - \frac{p}{c} + 1$. Consider an arbitrary point (L, t, m) during this process (note that L is the same). What is the number T of elements of S_{c+1} that are in X_i? Note that at $(L, t_c + 1, 1)$, $T = c - m_c$. Next, each time t increments, T increases by at least c (not counting the possible action inside S_{c+1}). Thus, $T \geq c - m_c + (t - t_c) \cdot c + m \geq (t - t_c) \cdot c + m + 1$. This helps us estimate the sum s_{last} of the elements of S_{c+1} in X_i at time (L, t, m). Note that the length of S_{c+1} is $L - t + 1$, so

$$s_{last} \geq \sum_{i=1}^{T+t-1} s_{c+1,i} = \alpha^{c-1} \cdot \sum_{i=1}^{T+t-1} s_{1,i} = \alpha^{c-1} \cdot \frac{\alpha^{T+t} - \alpha}{\alpha - 1}.$$

Now, we can bound the average

$$\begin{aligned} p \cdot \operatorname{avg} X_i &= \sum_{i=1}^{c} sum_i + s_{last} \geq \frac{\alpha^{c+1} - \alpha}{(\alpha - 1)^2} + \alpha^{c-1} \cdot \frac{\alpha^{T+t} - \alpha}{\alpha - 1} \\ &= \frac{1}{(\alpha - 1)^2} \cdot (\alpha^{T+t+c-1}(\alpha - 1) + \alpha^c - \alpha) \\ &\geq \frac{1}{(\alpha - 1)^2} \cdot (\alpha^m \cdot \alpha^{c(t-t_c+1)}(\alpha - 1) + \alpha^c - \alpha), \end{aligned}$$

and the max-over-average ratio

$$\begin{aligned} \frac{max_m}{\operatorname{avg} X_i} &\leq \frac{p(\alpha - 1)\alpha^m \cdot \alpha^{2tp/L} \cdot (1 - \alpha^{-2p/L})}{\alpha^m \cdot \alpha^{c(t-t_c+1)}(\alpha - 1) + \alpha^c - \alpha} \leq \frac{\ln 2 \cdot \alpha^{c+1} \cdot \alpha^{2tp/L} \cdot (1 - \alpha^{-2p/L})}{\alpha^{c-1} \cdot \alpha^{c(t-t_c+1)}(\alpha - 1) + \alpha^c - \alpha} \\ &= \frac{\ln 2 \cdot (1 - \alpha^{-2p/L}) \cdot \alpha^{2tp/L}}{\alpha^{c+2}(\alpha - 1) \cdot \alpha^{ct}\alpha^{-ct_c} + \alpha^{-1} - \alpha^{-c}} \end{aligned}$$

where we let m take its maximum possible value $m = c + 1$, and used the approximation $p(\alpha - 1) \approx \ln 2$. The last expression is a decreasing function of t, which can be checked by differentiation (the derivative is $A(2/L - c/p) + B$ for positive values A, B), so $t = t_c$ gives the worst-case bound

$$\frac{max_m}{\operatorname{avg} X_i} \leq \frac{\ln 2 \cdot (1 - 2^{-2/L}) \cdot 2^{2t_c/L}}{2^{\frac{c+2}{p}}(2^{1/p} - 1) + 2^{-1/p} - 2^{-c/p}} \leq \frac{4 \ln 2 \cdot (2^{2/L} - 1) \cdot 2^{-\frac{2p}{cL}}}{2^{\frac{c+2}{p}}(2^{1/p} - 1) + 2^{-1/p} - 2^{-c/p}},$$

where we also used the bound $t_c \leq L - p/c + 1$ (observed above) in the second inequality. This proves the third inequality of the lemma. □

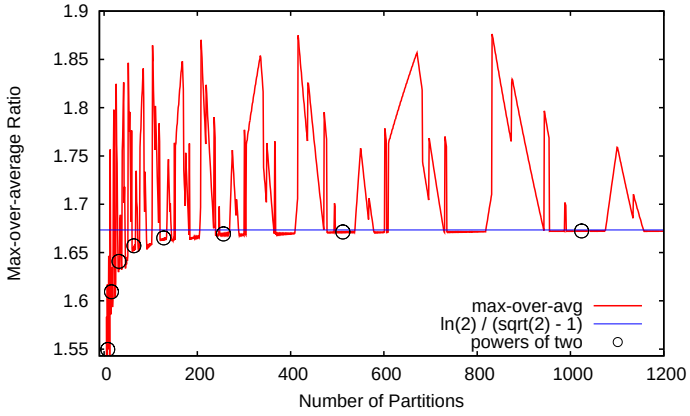

Figure 1: Extension of our scheme to arbitrary p. The encircled points correspond to the max-over-average ratio of powers of two. The blue curve is the bound proved in our analysis.

Combining the previous results, we obtain the following results:

Theorem 8. For every X_i in our scheme, the max-over-average ratio is bounded as:

$$\frac{\max X_i}{\operatorname{avg} X_i} \leq \frac{\ln 2}{\sqrt{2}-1} + O(1/p) \approx 1.673 + O(1/p) \ .$$

Proof. We can assume that p is large, e.g. $p \geq 256$, since we checked smaller values with a computer (see below). With this assumption, we can use the bounds in Lemma 5. The first bound gives the required approximation ratio for small values of l, $1 \leq l \leq 8$, by checking each case separately. The second bound of the lemma gives the ratio for $8 \leq l \leq p/16$. The remaining cases are when $l > p/16$, i.e. $c \leq 16$. These cases can again be individually checked for each c, by using the third bound given by Lemma 5. □

Corollary 2. If p is a power of two, then there is an algorithm for Flow with competitive ratio $\frac{\ln 2}{\sqrt{2}-1} + O(1/p) \approx 1.673 + O(1/p)$.

Last, similar to Corollary 1, via rounding, an algorithm for Flow can be used for Part while incurring an error term that depends on the weights of the sequence.

Corollary 3. There is a deterministic preemptive online algorithm for Part with competitive ratio $1.68 + O(1/p) + O(max/S)$, where S is the total weight of the sequence and max is the maximum element of the sequence.

Extending Our Scheme to Arbitrary Number of Partitions. Our scheme can equally be applied to arbitrary values of p, and we experimentally verified that max-over-average ratios better than two are obtained. Conducting a rigorous analysis for values of p that are not a power of two poses further complications and proves challenging. A case in point are the maxima of super-partitions, which could not be computed according to Lemma 4, and many rounding problems.

In Figure 1, we plot the max-over-average ratio of our scheme against the number of partitions p. The plot shows that the power of two cases are the ones that give the best competitive ratio. The chaotic behavior of the plot also indicates that many mechanisms within our scheme are at work at the same time that need to be bounded appropriately. We thus pose as an open problem: Can we obtain an analyzable partitioning scheme for arbitrary p with max-over-average ratio strictly better than 2?

5.2.2 Lower Bound

Consider an arbitrary current partitioning in Flow, and let $x_{\max}$ denote the current bottleneck value. For an appropriately chosen value $\alpha \in (1, 2)$, partition the current partitions into two categories L, H, so

that partitions in L are of weight at most $\frac{x_{\max}}{\alpha}$, and H are all other partitions. Our argument is based on the observation that if too many partition are in L, then the current competitive ratio cannot be good. However, if only very few partitions are in L, then after consuming only very little flow from the input, two partitions in H need to be merged with each other, which results in a large increase in the bottleneck value. This idea is formalized in the following theorem:

Theorem 9. Every deterministic preemptive online algorithm for FLOW has an approximation factor of at least $\frac{5}{2} - \sqrt{2} \approx 1.086$.

Proof. Let ALG be a deterministic f-approximation algorithm, for some $f > 1$. Suppose that the input is a sequence of unit weight requests. Consider a point in time when the total weight $x_{\max}$ of a heaviest partition is large enough so that rounding effects do not matter. Denote the weights of the partitions by $a_1, \ldots, a_p$. Since ALG is an f-approximation algorithm,

$$x_{\max} \leq \frac{f}{p} \sum_{i=1}^{p} a_i. \tag{5}$$

Let $\alpha \in (1, 2)$ be a parameter whose value will be optimized later. Partition the set of indices $[p]$ into sets L and H such that $i \in L$ if $a_i < \frac{x_{\max}}{\alpha}$, and $i \in H$ otherwise. Using this definition in Inequality 5, we obtain

$$\frac{p x_{\max}}{f} \quad \leq \quad |L| \cdot \frac{x_{\max}}{\alpha} + (p - |L|) \cdot x_{\max} \qquad \Leftrightarrow \qquad |L| \leq p \cdot \frac{1 - 1/f}{1 - 1/\alpha}.$$

We will prove now that after processing the next $(|L|+1)\frac{2x_{\max}}{\alpha}$ requests from the input sequence, there is at least one partition of weight at least $\frac{2x_{\max}}{\alpha}$. Indeed, consider the sequence:

$$X = a_1 a_2 \ldots a_p \quad \underbrace{1 \ldots 1}_{(|L|+1)\frac{2x_{\max}}{\alpha}}$$

and, for the sake of a contradiction, suppose that there is a partitioning of this sequence into p parts with maximum partition weight smaller than $\frac{2x_{\max}}{\alpha}$. By Lemma 1, such a partitioning would be found by a run of PROBE($\frac{2x_{\max}}{\alpha} - 1$) on X read from right to left. Notice first that every resulting partition contains at most one element from H (call such an element heavy), as the sum of two heavy elements already exceeds the bound $\frac{2x_{\max}}{\alpha} - 1$. The first $L + 1$ partitions created by PROBE each consist of 1s only. Since there are $|H| = p - |L|$ heavy elements, this implies that PROBE creates another $|H|$ partitions each containing one heavy element, which totals to $p + 1$ partitions, a contradiction.

Thus, after processing the next $(|L| + 1)\frac{2x_{\max}}{\alpha}$ ones from the input sequence, the average weight of a partition is bounded by:

$$\begin{aligned} A \quad &\leq \quad \frac{|H| \cdot x_{\max} + |L| \cdot \frac{x_{\max}}{\alpha} + (|L|+1) \cdot \frac{2x_{\max}}{\alpha}}{p} = x_{\max} \cdot \frac{(p - |L|) + |L| \cdot \frac{3}{\alpha} + \frac{2}{\alpha}}{p} \\ &= \quad x_{\max} \left(1 + |L| \frac{\frac{3}{\alpha} - 1}{p} + \frac{2}{p\alpha}\right) \leq x_{\max} \left(1 + p \cdot \frac{1 - \frac{1}{f}}{1 - \frac{1}{\alpha}} \cdot \frac{\frac{3}{\alpha} - 1}{p} + \frac{2}{p\alpha}\right) \\ &= \quad x_{\max} \left(\frac{2f + \alpha - 3}{f(\alpha - 1)} + \frac{2}{p\alpha}\right), \end{aligned}$$

where we applied Inequality 6, and as proved above, the maximum value then is at least $x'_{\max} \geq \frac{2x_{\max}}{\alpha}$. Since ALG is a f-approximation algorithm, we have $x'_{\max} \leq f \cdot A$, so

$$\frac{2x_{\max}}{\alpha} \leq f \cdot x_{\max} \left(\frac{2f + \alpha - 3}{f(\alpha - 1)} + \mathrm{O}(1/p)\right),$$

which implies that $2f \geq 5 - \left(\alpha + \frac{2}{\alpha}\right) - \mathrm{O}(1/p)$. The function $g(\alpha) = \alpha + 2/\alpha$ achieves its minimum at $\alpha = \sqrt{2}$, which implies that $f \geq \frac{5}{2} - \sqrt{2} - \mathrm{O}(1/p)$, giving the result. □

This lower bound argument shows that competitive ratios of at least 1.086 occur repeatedly while processing the input sequence. Recall that the lower bound for PART given in Theorem 7 only holds for two specific input lengths.

References

[1] S. H. Bokhari. Partitioning problems in parallel, pipeline, and distributed computing. *IEEE Trans. Comput.*, 37(1):48–57, January 1988. URL: http://dx.doi.org/10.1109/12.75137, doi:10.1109/12.75137.

[2] Allan Borodin and Ran El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, New York, NY, USA, 1998.

[3] Ashish Chiplunkar, Sumedh Tirodkar, and Sundar Vishwanathan. On randomized algorithms for matching in the online preemptive model. In *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pages 325–336, 2015.

[4] Yuval Emek, Pierre Fraigniaud, Amos Korman, and Adi Rosén. Online computation with advice. *Theor. Comput. Sci.*, 412(24):2642–2656, May 2011.

[5] Leah Epstein, Asaf Levin, Danny Segev, and Oren Weimann. Improved bounds for online preemptive matching. In *30th International Symposium on Theoretical Aspects of Computer Science, STACS 2013, February 27 - March 2, 2013, Kiel, Germany*, pages 389–399, 2013.

[6] Joan Feigenbaum, Sampath Kannan, Andrew McGregor, Siddharth Suri, and Jian Zhang. On graph problems in a semi-streaming model. *Theor. Comput. Sci.*, 348(2):207–216, December 2005. URL: http://dx.doi.org/10.1016/j.tcs.2005.09.013, doi:10.1016/j.tcs.2005.09.013.

[7] Greg N. Frederickson. Optimal algorithms for tree partitioning. SODA '91, pages 168–177, Philadelphia, PA, USA, 1991. URL: http://dl.acm.org/citation.cfm?id=127787.127822.

[8] Edward F. Grove. Online bin packing with lookahead. In *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '95, pages 430–436, 1995.

[9] Magnús M. Halldórsson, Kazuo Iwama, Shuichi Miyazaki, and Shiro Taketomi. Online independent sets. *Theor. Comput. Sci.*, 289(2):953–962, October 2002.

[10] Yijie Han, Bhagirath Narahari, and Hyeong-Ah Choi. Mapping a chain task to chained processors. *Inf. Process. Lett.*, 44(3):141–148, 1992. URL: http://dblp.uni-trier.de/db/journals/ipl/ipl44.html#HanNC92.

[11] Pierre Hansen and Keh-Wei Lih. Improved algorithms for partitioning problems in parallel, pipelined, and distributed computing. *IEEE Trans. Comput.*, 1992.

[12] Chinmay Karande, Aranyak Mehta, and Pushkar Tripathi. Online bipartite matching with unknown distributions. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 587–596, New York, NY, USA, 2011. ACM. URL: http://doi.acm.org/10.1145/1993636.1993715, doi:10.1145/1993636.1993715.

[13] R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the Twenty-second Annual ACM Symposium on Theory of Computing*, STOC '90, pages 352–358, New York, NY, USA, 1990. ACM. URL: http://doi.acm.org/10.1145/100216.100262, doi:10.1145/100216.100262.

[14] Sanjeev Khanna, S. Muthukrishnan, and Steven Skiena. Efficient array partitioning. In *ICALP*, volume 1256, pages 616–626. Springer Berlin Heidelberg, 1997. URL: http://dx.doi.org/10.1007/3-540-63165-8_216.

[15] Christian Konrad. Two-constraint domain decomposition with space filling curves. *Parallel Comput.*, 37(4-5):203–216, April 2011. URL: http://dx.doi.org/10.1016/j.parco.2011.03.002, doi:10.1016/j.parco.2011.03.002.

[16] Christian Konrad. Streaming partitioning of sequences and trees. In *19th International Conference on Database Theory, ICDT 2016, Bordeaux, France, March 15-18, 2016*, pages 13:1–13:18, 2016.

[17] Mark S. Manasse, Lyle A. McGeoch, and Daniel D. Sleator. Competitive algorithms for server problems. *J. Algorithms*, 11(2):208–230, May 1990. URL: http://dx.doi.org/10.1016/0196-6774(90)90003-W, doi:10.1016/0196-6774(90)90003-W.

[18] Fredrik Manne and Bjørn Olstad. Efficient partitioning of sequences. *IEEE Trans. Comput.*, 44(11):1322–1326, November 1995. URL: http://dx.doi.org/10.1109/12.475128, doi:10.1109/12.475128.

[19] Fredrik Manne and Tor Sørevik. Optimal partitioning of sequences. *J. Algorithms*, 19(2):235–249, September 1995. URL: http://dx.doi.org/10.1006/jagm.1995.1035, doi:10.1006/jagm.1995.1035.

[20] Serge Miguet and Jean-Marc Pierson. Heuristics for 1d rectilinear partitioning as a low cost and high quality answer to dynamic load balancing. HPCN Europe '97, pages 550–564, London, UK, UK, 1997. Springer-Verlag. URL: http://dl.acm.org/citation.cfm?id=645561.659355.

[21] Ali Pinar and Cevdet Aykanat. Fast optimal load balancing algorithms for 1d partitioning. *J. Parallel Distrib. Comput.*, 64(8):974–996, August 2004. URL: http://dx.doi.org/10.1016/j.jpdc.2004.05.003, doi:10.1016/j.jpdc.2004.05.003.

[22] Isabelle Stanton. Streaming balanced graph partitioning algorithms for random graphs. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, pages 1287–1301. SIAM, 2014. URL: http://dl.acm.org/citation.cfm?id=2634074.2634169.

[23] Isabelle Stanton and Gabriel Kliot. Streaming graph partitioning for large distributed graphs. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '12, pages 1222–1230, New York, NY, USA, 2012. ACM. URL: http://doi.acm.org/10.1145/2339530.2339722, doi:10.1145/2339530.2339722.

[24] J. D. Ullman. The performance of a memory allocation algorithm. Technical report, Princeton University, 1971.

[25] Jeffrey S. Vitter. Random sampling with a reservoir. *ACM Trans. Math. Softw.*, 11(1):37–57, March 1985. URL: http://doi.acm.org/10.1145/3147.3165, doi:10.1145/3147.3165.

[26] Andrew Chi-Chin Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, SFCS '77, pages 222–227, Washington, DC, USA, 1977. IEEE Computer Society. URL: http://dx.doi.org/10.1109/SFCS.1977.24, doi:10.1109/SFCS.1977.24.

[27] Mariano Zelke. Weighted matching in the semi-streaming model. *Algorithmica*, 62(1-2):1–20, February 2012. URL: http://dx.doi.org/10.1007/s00453-010-9438-5, doi:10.1007/s00453-010-9438-5.

A Barely-random Algorithm for Unit Requests and $p = 2$

We show that using a single random bit gives a significant improvement over deterministic algorithms. The algorithm $\mathcal{A}_0$ places the separator at positions 2^i for odd or even i depending on an initial random choice.

Algorithm 5 $\mathcal{A}_0$

```
i ← 0 or 1 with probability 1/2 each
for each request j = 1 ... n do
  if j = 2^i then
    move the separator to the current position
    i → i + 2
  end if
end for
```

Denote the competitive ratio of $\mathcal{A}_0$ on sequences of length n by $R^n_{\mathcal{A}_0}$. Then, we obtain the following theorem:

Theorem 19. $\mathbb{E}[R^n_{\mathcal{A}_0}] = 1.5$.

Proof. Let $\alpha \in [0, 1)$ and $i \in \mathbb{N}$ be such that the sequence length is $n = 2 \cdot 2^{i+\alpha}$. The bottleneck value of an optimal partition is at least $\lceil \frac{n}{2} \rceil = \lceil 2^{i+\alpha} \rceil$. Since, $2^{i+2} > n$, the algorithm puts the separator either at position 2^i or 2^{i+1}, each with probability $1/2$. In the first case, the bottleneck value is $n - 2^i$, while in the second case it is 2^{i+1}. Thus,

$$E[R^n_{\mathcal{A}_0}] = \frac{1}{2} \cdot \frac{n - 2^i + 2^{i+1}}{\lceil n/2 \rceil} \leq 1 + \frac{1}{2^{\alpha+1}} \leq 1.5 \ .$$

□

B Complete Scheme for $p = 4$

$$\overbrace{\boldsymbol{x^{\frac{1}{4}}, x^{\frac{2}{4}}, x^{\frac{3}{4}}, x^{\frac{4}{4}}}}^{S_1} \mid \overbrace{x^{\frac{2}{4}}, x^{\frac{3}{4}}, x^{\frac{4}{4}}, x^{\frac{5}{4}}}^{S_2} \mid \overbrace{x^{\frac{3}{4}}, x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}}}^{S_3} \mid \overbrace{x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}}, x^{\frac{7}{4}}}^{S_4}$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{3}{4}}, x^{\frac{4}{4}}} \mid \boldsymbol{x^{\frac{2}{4}}}, x^{\frac{3}{4}}, x^{\frac{4}{4}}, x^{\frac{5}{4}} \mid x^{\frac{3}{4}}, x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}} \mid x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}}, x^{\frac{7}{4}}$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{3}{4}}, x^{\frac{4}{4}}} \mid \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}})}, x^{\frac{4}{4}}, x^{\frac{5}{4}} \mid x^{\frac{3}{4}}, x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}} \mid x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}}, x^{\frac{7}{4}}$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{3}{4}}, x^{\frac{4}{4}}} \mid \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}})}, x^{\frac{4}{4}}, x^{\frac{5}{4}} \mid x^{\frac{3}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{5}{4}}, x^{\frac{6}{4}} \mid x^{\frac{4}{4}}, x^{\frac{5}{4}}, x^{\frac{6}{4}}, x^{\frac{7}{4}}$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{3}{4}}, x^{\frac{4}{4}}} \mid \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}})}, x^{\frac{4}{4}}, x^{\frac{5}{4}} \mid x^{\frac{3}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{5}{4}}, x^{\frac{6}{4}} \mid x^{\frac{4}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{6}{4}}, x^{\frac{7}{4}}$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{3}{4}}(1 + x^{\frac{1}{4}})} \mid \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{4}{4}}}, x^{\frac{5}{4}} \mid x^{\frac{3}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{5}{4}}, x^{\frac{6}{4}} \mid x^{\frac{4}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{6}{4}}, x^{\frac{7}{4}}$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{3}{4}}(1 + x^{\frac{1}{4}})} \mid \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{4}{4}}(1 + x^{\frac{1}{4}})} \mid x^{\frac{3}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{5}{4}}, x^{\frac{6}{4}} \mid x^{\frac{4}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{6}{4}}, x^{\frac{7}{4}}$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{3}{4}}(1 + x^{\frac{1}{4}})} \mid \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{4}{4}}(1 + x^{\frac{1}{4}})} \mid x^{\frac{3}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{5}{4}}, x^{\frac{6}{4}} \mid x^{\frac{4}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{6}{4}}, x^{\frac{7}{4}}$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{3}{4}}(1 + x^{\frac{1}{4}})} \mid \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{4}{4}}(1 + x^{\frac{1}{4}})} \mid x^{\frac{3}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{5}{4}}(1 + x^{\frac{1}{4}}) \mid x^{\frac{4}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{6}{4}}, x^{\frac{7}{4}}$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{3}{4}}(1 + x^{\frac{1}{4}})} \mid \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{4}{4}}(1 + x^{\frac{1}{4}})} \mid x^{\frac{3}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{5}{4}}(1 + x^{\frac{1}{4}}) \mid x^{\frac{4}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{6}{4}}(1 + x^{\frac{1}{4}})$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}})(1 + x^{\frac{2}{4}})} \mid \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{4}{4}}(1 + x^{\frac{1}{4}})} \mid \boldsymbol{x^{\frac{3}{4}}(1 + x^{\frac{1}{4}})}, x^{\frac{5}{4}}(1 + x^{\frac{1}{4}}) \mid x^{\frac{4}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{6}{4}}(1 + x^{\frac{1}{4}})$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}})(1 + x^{\frac{2}{4}})} \mid \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}})(1 + x^{\frac{2}{4}})} \mid \boldsymbol{x^{\frac{3}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{5}{4}}(1 + x^{\frac{1}{4}})} \mid x^{\frac{4}{4}}(1 + x^{\frac{1}{4}}), x^{\frac{6}{4}}(1 + x^{\frac{1}{4}})$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}})(1 + x^{\frac{2}{4}})} \mid \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}})(1 + x^{\frac{2}{4}})} \mid \boldsymbol{x^{\frac{3}{4}}(1 + x^{\frac{1}{4}})(1 + x^{\frac{2}{4}})} \mid \boldsymbol{x^{\frac{4}{4}}(1 + x^{\frac{1}{4}})}, x^{\frac{6}{4}}(1 + x^{\frac{1}{4}})$$

$$\boldsymbol{x^{\frac{1}{4}}(1 + x^{\frac{1}{4}})(1 + x^{\frac{2}{4}})} \mid \boldsymbol{x^{\frac{2}{4}}(1 + x^{\frac{1}{4}})(1 + x^{\frac{2}{4}})} \mid \boldsymbol{x^{\frac{3}{4}}(1 + x^{\frac{1}{4}})(1 + x^{\frac{2}{4}})} \mid \boldsymbol{x^{\frac{4}{4}}(1 + x^{\frac{1}{4}})(1 + x^{\frac{2}{4}})}$$